

A6826A – PCI-X Dual port 2Gb/s Fibre Channel Adapter

Performance Paper for PCI platforms



Table of contents

Introduction	2
Executive summary.....	2
Test results	3
IOPs	3
Service demand.....	3
Throughput.....	4
Scalability	5
System configuration guidelines.....	7
Test details.....	8
Products used in testing.....	8
Test configuration.....	9
Additional Information.....	9

Introduction

This paper provides basic I/O performance and scalability information for the A6826A dual-channel 2Gb Fibre Channel adapter on HP's PCI platforms. A series of tests were conducted to evaluate the adapter's performance on HP servers, namely rp5450, rp7400, rp8400 and Superdome. The tests measured the IOPS and throughput of single and multiple adapters.

This paper addresses the performance capabilities of A6826A when used with the above mentioned HP platforms. The paper also discusses the system setup considerations to obtain the maximum possible performance from A6826A on these platforms.

This paper focuses on the following topics:

- **Test results:** Single and multiple adapters performance data, which include the IOPS and throughput for single and dual ports, will be discussed.
- **Scalability:** Multiple adapter scalability tests on systems namely Superdome, rp8400, rp7400 and rp5450 will be discussed.
- **System configuration guidelines:** Superdome, rp8400, rp7400 and rp5450 system configurations and recommendations will be discussed.
- **Test details:** HP's products used, test setup, benchmark tool used and system configuration in the test setup will be discussed.

Executive summary

A6826A on HP mid range and high end platforms offers an excellent SAN solution. A6826A provides read throughput of 195MB/s, and write throughput of 188MB/s on a single port. With dual ports a read throughput of 390MB/s, and write throughput of 368MB/s are achieved.

A6826A performance summary

Operation	single port			two ports		
	read	write	bd*	read	write	bd
iops	35530	21300	20300	70100	42700	31300
throughput (MB/s)	achieved	195	188	351	390	368
	% of 2Gb/s FC theoretical max	97.5	94	87.7	97.5	92
						51.1

NOTE:

*bd: Bidirectional Operation.

**The full duplex throughput is limited by 66MHz PCI bandwidth (528 MB/s).

The above performance data is obtained on A6826A seated in a 66MHz PCI slot of a single cell single CPU (750MHz) rp8400.

On rp8400 and Superdome A6826A provides outstanding performance with linear scaling up to 4 cards, when all the eight ports are being used at 2Gb/s bandwidth. Additional cards may be added to provide greater connectivity. The card delivers excellent performance in limited configurations on rp7400 and rp5450.

Test results

Diskbench (db) utility is used to generate the read and write test traffic. Various block size tests were executed for read and write on single and dual ports. The IOPS metric are obtained with 1 KB block size transfers. The throughput metric is obtained with 128 KB block size transfers. The throughput metric is useful in modeling large sequential transfers such as remote backup etc. The IOPS metric is useful in modeling small transactional traffic.

IOPs

Chart 1a: IOPS (in thousands)

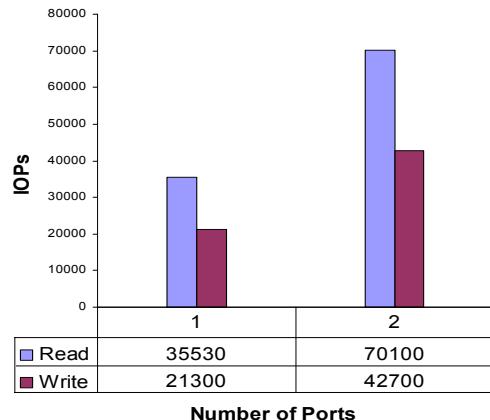
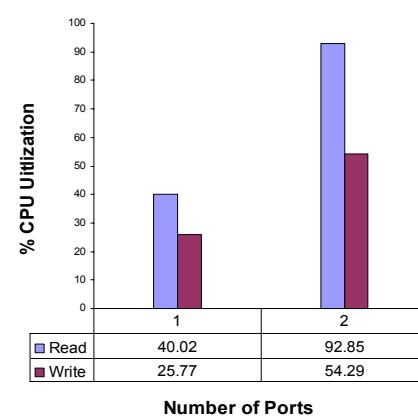


Chart 1b: % CPU Utilization



The chart, chart 1a, shows the number of IO operations per second for read and write operations on single and dual ports of A6826A. The X-axis is the number of ports and the Y-axis the number of IO operations per second.

The chart 1a shows linear scaling of IOPS for single and dual ports. The number of read IO operations per second for single port is 35530 and for dual ports are 70100. The number of write operations per second for single port is 21300 and for dual port are 42700. The IOPS metric for A6826A is limited by the processor used on A6826A.

The chart 1b shows the CPU utilization for the IO operation tests.

The chart 1b shows the % CPU utilization for IO operation tests. The X-axis is number of the Fibre Channel ports. The Y-axis is the % CPU utilization. Only one 1 CPU is configured in these tests, hence this is a single CPU utilization.

Service demand

A6826A offers a great service demand for small size IO operations. To illustrate service demand of A6826A small size operations tests were conducted using db. Read and write operations of IO sizes of 4KB and 8KB were performed on a single cell single CPU rp8400 and the results were recorded.

The following table shows single and dual port throughput and CPU utilization for Sequential Read and Sequential Writes with 4KB and 8KB IO sizes.

Number of Ports	Sequential Read						Sequential Write					
	4KB			8KB			4KB			8KB		
	Thru ¹	CPU ²	SD ³	Thru	CPU	SD	Thru	CPU	SD	Thru	CPU	SD
Single	130	37%	2.93	190	25%	1.34	83	25%	3.08	127	20%	1.61
Dual	275	74%	2.75	379	65%	1.75	164	53%	3.30	248	58%	2.39

NOTE:

¹ Throughput in MB/s

² Single CPU utilization

³ Service Demand = $((\% \text{CPU utilization} / 100) / \text{Throughput in KB/sec})$

Data gathered on single cell single CPU rp8400

Throughput

Chart 2a : Single and Dual Port Throughput

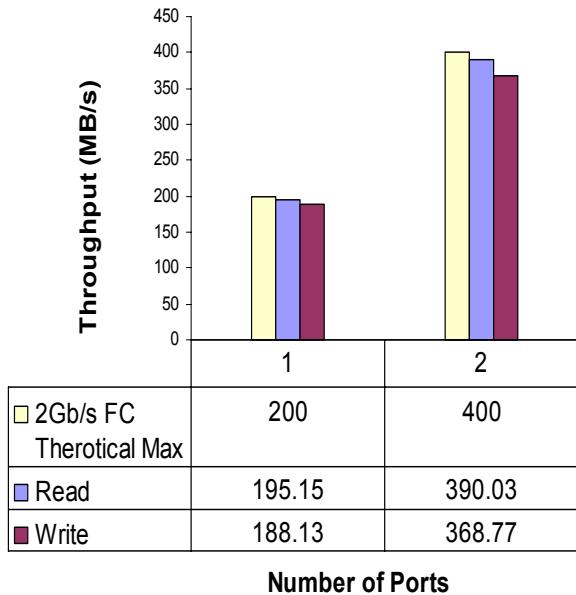
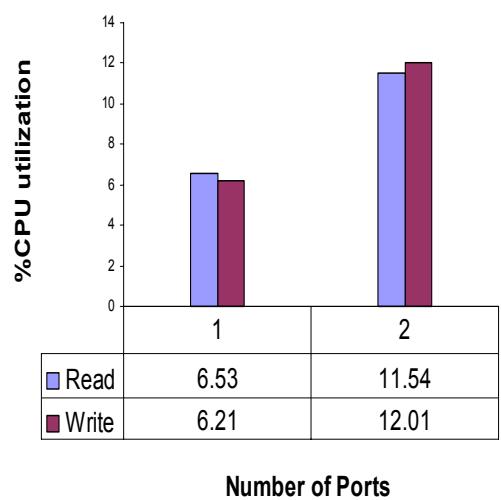


Chart 2b : CPU Utilization for Throughput tests



Data obtained on a single cell, single CPU, rp8400

The chart 2a shows the read and write throughput of single and dual ports of A6826A. The X-axis is the number of Fibre Channel ports. The Y-axis is the throughput in MB/s.

As the chart 2a shows, one port of A6828A performs read at 195MB/s. Two ports scale at 2x the performance of single port performing reads up to 383MB/s.

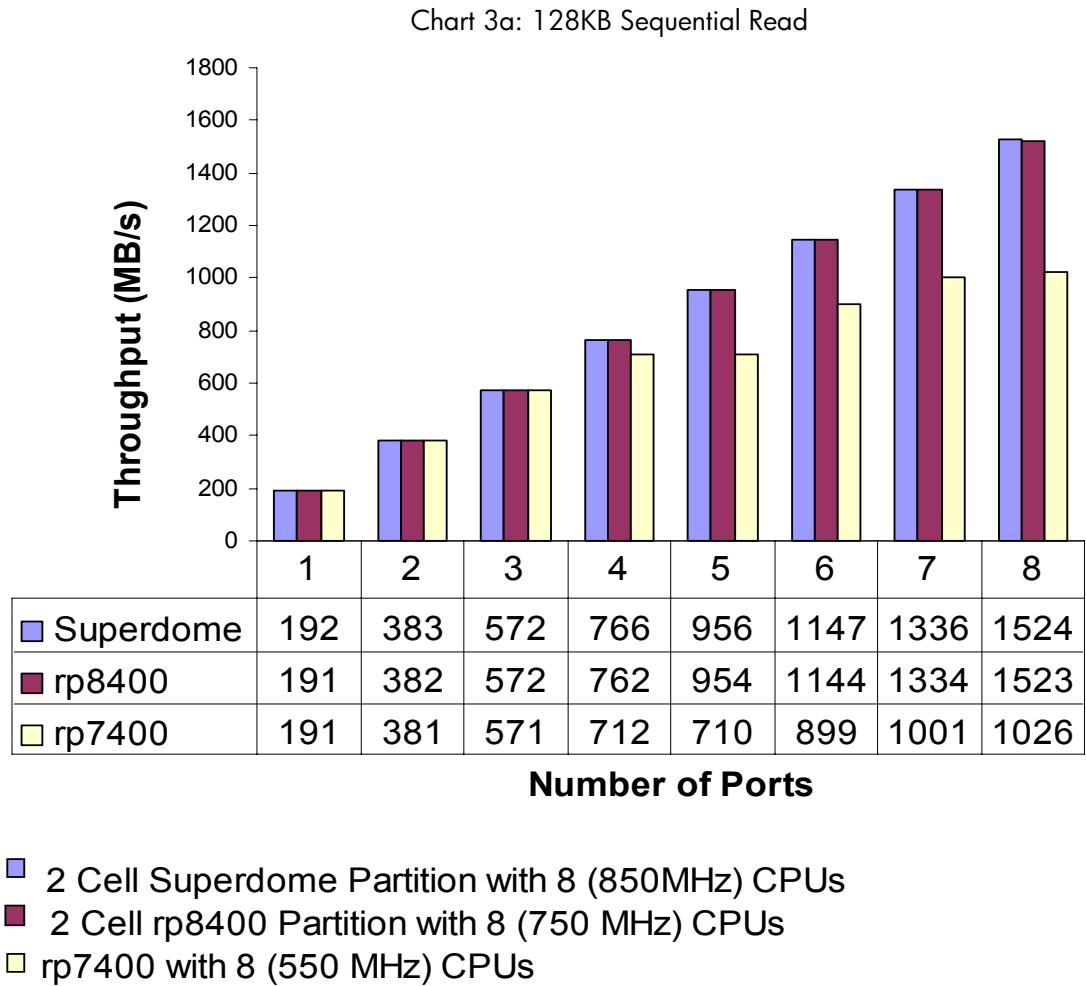
For the write operations, one port of A6826A performs write operations at 188MB/s. Two ports scale approximately at 2x the performance of the single port performing writes up to 368MB/s.

The chart 2a demonstrates outstanding read and write performance with linear scaling for 1 and 2 ports of A6826A.

The chart 2b shows the % CPU utilization of throughput tests. The X-axis is the number of the Fibre Channel port. The Y-axis is the % CPU utilization. Only one 1 CPU is configured in these tests, hence this is a single CPU utilization. The charts 2b shows very small percentages of CPU utilization for single and dual ports in read and write throughput tests.

Scalability

A series of tests were conducted to evaluate the scalability aspect of A6826A on various HP systems, namely Superdome, rp8400, rp7400 and rp5450.

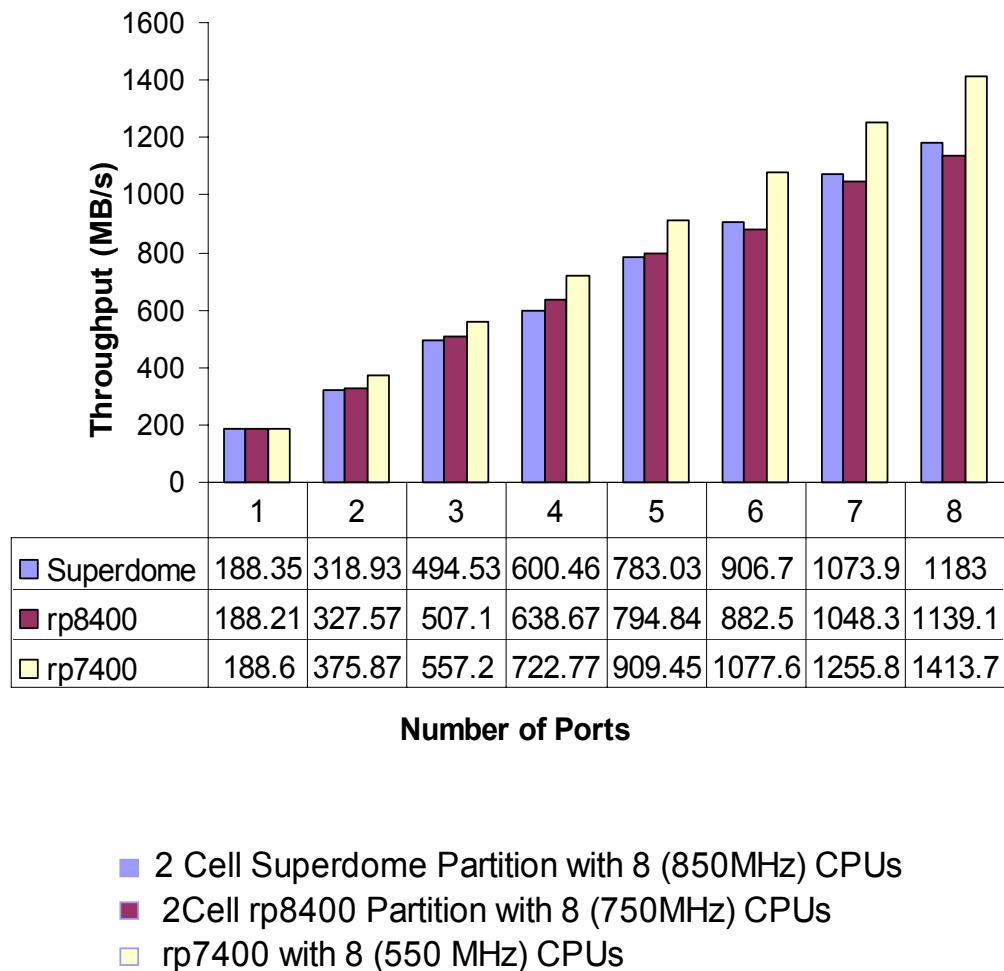


The above chart, chart 3a, shows the read throughput performance of 1, 2, 3 and 4 A6826A in Superdome, rp8400 and rp7400. The X-axis is the number of ports. Each port represents one port of A6826A. The Y-axis is the throughput performance in MB/s.

The chart shows excellent read throughput performance with linear scaling for 4 A6826As on Superdome and rp8400 reaching an aggregate of 1524MB/s. Superdome and rp8400 along with A6826A offer a great scalable and connectivity SAN solution.

On rp7400 the read throughput scales linearly up to 4 ports (2 A6826As) and reaches a plateau from 4 to 5 ports and then ascends up to 1025MB/s for 8 ports. The decline in throughput from 7 to 8 results from the IO subsystem limitations on rp7400. rp7400 offers a good scalable and connectivity SAN solution.

Chart 3b: 128 KB Sequential Write



The above chart, chart 3b, shows the write throughput performance of 1, 2, 3 and 4 A6826A in Superdome, rp8400, rp7400 and rp5450 systems. The X-axis is the number of ports. Each port represents one port of A6826A. The Y-axis is the throughput performance in MB/s.

The chart shows excellent write throughput performance with linear scaling for 4 A6826As on Superdome and rp8400 reaching an aggregate of 1180 MB/s. Superdome and rp8400 along with A6826A offer a great scalable and connectivity solution.

On rp7400 the write throughput scales linearly for 4 A6286As reaching an aggregate of 1413MB/s.

The table below summarizes the maximum number of cards that can be used on these platforms in order achieve linear scalability. Additional cards can be used to provide better connectivity.

Ports operating speed	Maximum number of A6826A for linear scalability			
	rp8400	Superdome	rp7400	rp5450
Both Ports at 2Gb	4	2 per IO Cage	4	2
One port at 1 Gb and other at 2Gb	6	3 per IO Cage	6	4
Both ports at 1 Gb	8	4 per IO Cage	8	4

System configuration guidelines

Superdome

To get the throughput and scalability with multiple cards on a Superdome, HP recommends the following:

The cells should be configured with a multiple of 8 equal capacity DIMMs to take advantage of memory interleaving. The 8 DIMMs should be evenly distributed across the two busses.

Each IO cage of Superdome offers twelve PCI slots, eight 2X and four 4X PCI slots. The 2X slots have 265MB/s bandwidth and 4X PCI slots have 530 MB/s bandwidth. HP recommends A6826A to be installed in a 4X PCI slot (Physical slot number 4, 5, 6 or 7) in order to achieve performance shown in this paper. To get a linear scalability HP recommends up to a maximum of two A6826As per IO cage.

rp8400

To get the throughput and scalability with multiple cards on an rp8400, HP recommends the following:

The cells should be configured with a multiple of 8 equal capacity DIMMs to take advantage of memory interleaving. The 8 DIMMs should be evenly distributed across the two busses.

rp8400's IO subsystem offers sixteen PCI slots, two 2X PCI and fourteen 4X PCI slots. The 2X slots have 265MB/s bandwidth and 4X PCI slots have 530 MB/s bandwidth. HP recommends A6826A to be installed in a 4X PCI slot (Physical slot numbers 1, 2, 3, 4, 5, 6) to achieve performance shown in this paper. To get a linear scalability HP recommends a rp8400 should have at least 2 cell partition and with two A6826As installed in any of 1 to 8 4X PCI slots and a maximum of two A6826As in any of the 9 to 16 4X PCI slots.

rp7400

To get the throughput and scalability with multiple cards on an rp7400, HP recommends the following:

The rp7400 supports up to 4 memory carriers each providing 8 memory slots. All the carriers should be populated with equal size DIMMs to maximize the memory bandwidth.

rp7400 offers twelve PCI slots, two 2X (a.k.a. Turbo) and ten 4X (a.k.a. Twin Turbo) PCI slots. HP recommends A6826A to be installed in a 4X slot (Physical slot numbers 3, 4, 5, 6, 7, 8, 9, 10, 11, 12) to achieve performance shown in this paper. To achieve a linear scalability HP recommends up to a maximum of two A6826As to be installed in the right IO backplane (Physical slots 7 to 12) and up to a maximum of two A6826As to be installed in the left IO backplane (Physical slots 1 to 6).

rp5450

rp5450 is designed to be low end mid range system with limited IO capabilities. HP recommends A6826A on rp5450 to be used as a connectivity solution.

Test details

The performance results presented in this paper were obtained with A6826A on various HP platforms. The system configurations in the test setup are tabulated in the following table:

Products used in testing

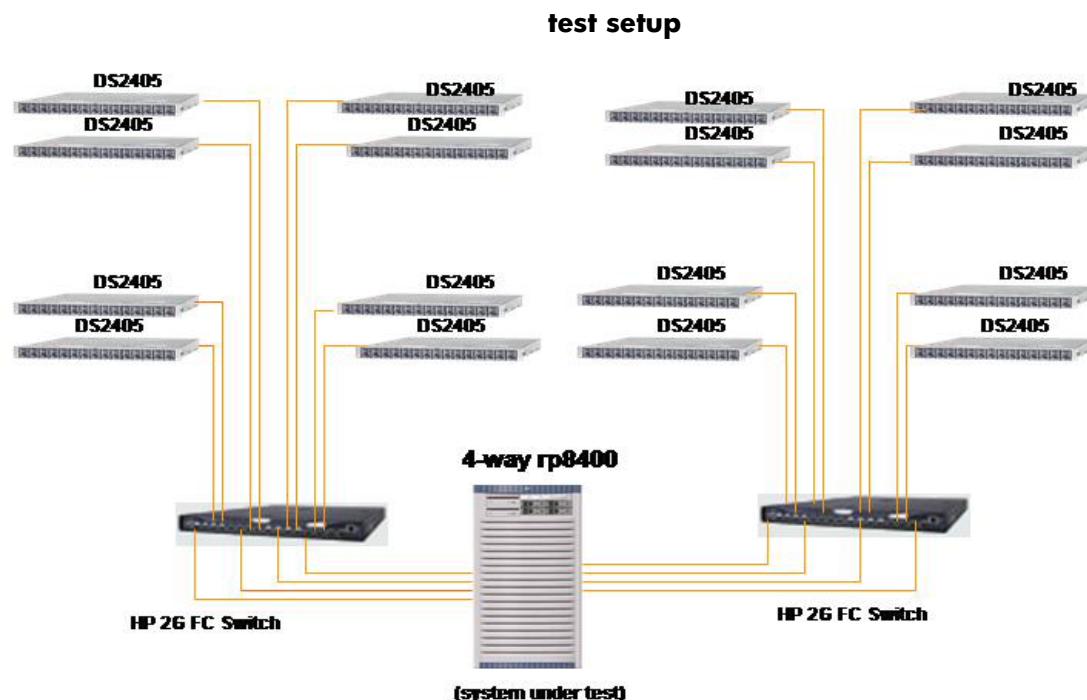
products used diskbench for performance measurement test		
Servers Tested		
		rp8400 <ul style="list-style-type: none"> • 2 cell partition • Each cell with <ul style="list-style-type: none"> ➢ 2 - 750MHz PA-8700 CPUs ➢ 4GB System Memory • HP-UX B.11.11 OS • 2Port 2Gig Fibre Channel driver (Ver. B.11.11.01)
		Superdome <ul style="list-style-type: none"> • 2 Cell partition • Each cell with <ul style="list-style-type: none"> ➢ 2 - 875 MHz PA-? CPUs ➢ 4 GB System Memory • HP-UX B.11.11 OS • 2Port 2Gig Fibre Channel driver (Ver. B.11.11.01)
		rp7400 <ul style="list-style-type: none"> • 8 – 550 MHz PA8600 CPUs • 8GB System Memory • HP-UX B.11.11 OS • 2Port 2Gig Fibre Channel driver (Ver. B.11.11.01)
		rp5450 <ul style="list-style-type: none"> • 4- 440 MHz PA8500 CPUs • 4GB System Memory • HP-UX b.11.11 OS • 2Port 2Gig Fibre Channel driver (Ver. B.11.11.01)
		A6826A HBA <ul style="list-style-type: none"> • PCI-X dual-channel 2Gb Fibre Channel adapter • Each port capable of independently operating at 1 or 2 Gb/s (Auto-Negotiation) • 33/66/100/133MHz-64bit
Benchmark software		Diskbench (db) is the benchmark suite that generated disk read and write traffic for these tests.
DS2405		HP StorageWorks 2Gb/s Disk System

Test configuration

The test configuration consists of 8-way (750 MHz) rp8400 with 8 GB of System memory and four A6826A HBAs. Two HBAs were connected to HP 2Gb/s Fibre Channel switches and the other two were connected to another HP 2Gb/s Fibre Channel switch. Sixteen HP DS2405 were evenly connected to the two switches. The switches were zoned so that each port of A6826A could see 2 DS2405.

Similar test setup was used to collect performance data on Superdome, rp8700 and rp5450.

Single CPU system configuration was to obtain IOPS and throughput results. Additional CPUs were added to study the scalability of A6826A on rp7400, rp8400 and Superdome.



Additional Information

For more information about A6826A please visit
<http://www.hp.com/products1/unixserverconnectivity/storagesnf2/index.html>.

© 2003 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license. Pentium is a U.S. registered trademark of Intel Corporation. Microsoft, Windows, and Windows NT are U.S. registered trademark of Microsoft Corporation. UNIX is a registered trademark of The Open Group.

07/2003

5981-7419EN

